

Theories of Vagueness

Rosanna Keefe
University of Sheffield



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK www.cup.cam.ac.uk
40 West 20th Street, New York, NY 10011–4211, USA www.cup.org
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

© Rosanna Keefe 2000

The book is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2000

Printed in the United Kingdom at the University Press, Cambridge

Typeset in Bembo 10.5/12pt System 3b2 [CE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Keefe, Rosanna.

Theories of vagueness / Rosanna Keefe.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 65067 4 (hardback)

1. Vagueness (Philosophy) I. Title.

B105.V33 K44 2000

160–dc21

ISBN 0 521 65067 4 hardback

Contents

<i>Acknowledgements</i>	<i>page xi</i>
Introduction	1
1 The phenomena of vagueness	6
2 How to theorise about vagueness	37
3 The epistemic view of vagueness	62
4 Between truth and falsity: many-valued logics	85
5 Vagueness by numbers	125
6 The pragmatic account of vagueness	139
7 Supervaluationism	152
8 Truth is super-truth	202
<i>References</i>	221
<i>Index</i>	229

The phenomena of vagueness

I. CENTRAL FEATURES OF VAGUE EXPRESSIONS

The parties to the vigorous debates about vagueness largely agree about which predicates are vague: paradigm cases include ‘tall’, ‘red’, ‘bald’, ‘heap’, ‘tadpole’ and ‘child’. Such predicates share three interrelated features that intuitively are closely bound up with their vagueness: they admit borderline cases, they lack (or at least apparently lack) sharp boundaries and they are susceptible to sorites paradoxes. I begin by describing these characteristics.

Borderline cases are cases where it is unclear whether or not the predicate applies. Some people are borderline tall: not clearly tall and not clearly not tall. Certain reddish-orange patches are borderline red. And during a creature’s transition from tadpole to frog, there will be stages at which it is a borderline case of a tadpole. To offer at this stage a more informative characterisation of borderline cases and the unclarity involved would sacrifice neutrality between various competing theories of vagueness. Nonetheless, when Tek is borderline tall, it does seem that the unclarity about whether he is tall is not merely epistemic (i.e. such that there is a fact of the matter, we just do not know it). For a start, no amount of further information about his exact height (and the heights of others) could help us decide whether he *is* tall. More controversially, it seems that there is no fact of the matter here about which we are ignorant: rather, it is *indeterminate* whether Tek is tall. And this indeterminacy is often thought to amount to the sentence ‘Tek is tall’ being neither true nor false, which violates the classical principle of bivalence. The law of excluded middle may also come into question when we consider instances such as ‘either Tek is tall or he is not tall’.

Second, vague predicates apparently lack well-defined extensions. On a scale of heights there appears to be no sharp boundary between

the tall people and the rest, nor is there an exact point at which our growing creature ceases to be a tadpole. More generally, if we imagine possible candidates for satisfying some vague F to be arranged with spatial closeness reflecting similarity, no sharp line can be drawn round the cases to which F applies. Instead, vague predicates are naturally described as having fuzzy, or blurred, boundaries. But according to classical logic and semantics all predicates have well-defined extensions: they cannot have fuzzy boundaries. So again this suggests that a departure from the classical conception is needed to accommodate vagueness.

Clearly, having fuzzy boundaries is closely related to having borderline cases. More specifically, it is the *possibility* of borderline cases that counts for vagueness and fuzzy boundaries, for if all actually borderline tall people were destroyed, ‘tall’ would still lack sharp boundaries. It might be argued that for there to be no sharp boundary between the F s and the not- F s just *is* for there to be a region of possible borderline cases of F (sometimes known as the penumbra). On the other hand, if the range of possible borderline cases between the F s and the not- F s was itself sharply bounded, then F would have a sharp boundary too, albeit one which was shared with the borderline F s, not with the things that were definitely not F . The thought that our vague predicates are not in fact like this – their borderline cases are not sharply bounded – is closely bound up with the key issue of higher-order vagueness, which will be discussed in more detail in §6.

Third, typically vague predicates are susceptible to sorites paradoxes. Intuitively, a hundredth of an inch cannot make a difference to whether or not a man counts as tall – such tiny variations, undetectable using the naked eye and everyday measuring instruments, are just too small to matter. This seems part of what it is for ‘tall’ to be a *vague* height term lacking sharp boundaries. So we have the principle $[S_1]$ if x is tall, and y is only a hundredth of an inch shorter than x , then y is also tall. But imagine a line of men, starting with someone seven feet tall, and each of the rest a hundredth of an inch shorter than the man in front of him. Repeated applications of $[S_1]$ as we move down the line imply that each man we encounter is tall, however far we continue. And this yields a conclusion which is clearly false, namely that a man less than five feet tall, reached after three thousand steps along the line, is also tall.

Similarly there is the ancient example of the heap (Greek *soros*, from which the paradox derives its name). Plausibly, [S₂] if x is a heap of sand, then the result y of removing one grain will still be a heap – recognising the vagueness of ‘heap’ seems to commit us to this principle. So take a heap and remove grains one by one; repeated applications of [S₂] imply absurdly that the solitary last grain is a heap. The paradox is supposedly owed to Eubulides, to whom the liar paradox is also attributed. (See Barnes 1982 and Burnyeat 1982 for detailed discussion of the role of the paradox in the ancient world.)

Arguments with a sorites structure are not mere curiosities: they feature, for example, in some familiar ethical ‘slippery slope’ arguments (see e.g. Walton 1992 and Williams 1995). Consider the principle [S₃] if it is wrong to kill something at time t after conception, then it would be wrong to kill it at time t minus one second. And suppose we agree that it is wrong to kill a baby nine months after conception. Repeated applications of [S₃] would lead to the conclusion that abortion even immediately after conception would be wrong. The need to assess this kind of practical argumentation increases the urgency of examining reasoning with vague predicates.

Wright (1975, p. 333) coined the phrase *tolerant* to describe predicates for which there is ‘a notion of degree of change too small to make any difference’ to their applicability. Take ‘[is] tall’ (for simplicity, in mentioning predicates I shall continue, in general, to omit the copula). This predicate will count as tolerant if, as [S₁] claims, a change of one hundredth of an inch never affects its applicability. A tolerant predicate must lack sharp boundaries; for if F has sharp boundaries, then a boundary-crossing change, however small, will always make a difference to whether F applies.¹ Moreover, a statement of the tolerance of F can characteristically serve as the inductive premise of a sorites paradox for F (as in the example of ‘tall’ again).

Russell provides one kind of argument that predicates of a given class are tolerant: if the application of a word (a colour predicate, for example) is paradigmatically based on unaided sense perception, it surely cannot be applicable to only one of an indiscriminable pair (1923, p. 87). So such ‘observational’ predicates will be tolerant with

¹ Note that throughout this book, when there is no potential for confusion I am casual about omitting quotation marks when natural language expressions are not involved, e.g. when talking about the predicate F or the sentence p & $\neg p$.

The phenomena of vagueness

respect to changes too small for us to detect. And Wright develops, in detail, arguments supporting the thesis that many of our predicates are tolerant (1975 and 1976). In particular, consideration of the role of ostension and memory in mastering the use of such predicates appears to undermine the idea that they have sharp boundaries which could not be shown by the teacher or remembered by the learner. Arguments of this kind are widely regarded as persuasive: I shall refer to them as ‘typical arguments for tolerance’. A theory of vagueness must address these arguments and establish what, if anything, they succeed in showing, and in particular whether they show that the inductive premise of the sorites paradox holds.

Considerations like Russell’s and Wright’s help explain why vague predicates are so common (whatever we say about the sorites premise). And they also seem to suggest that we *could* not operate with a language free of vagueness. They make it difficult to see vagueness as a merely optional or eliminable feature of language. This contrasts with the view of vagueness as a defect of natural languages found in Frege (1903, §56) and perhaps in Russell’s uncharitable suggestion (1923, p. 84) that language is vague because our ancestors were lazy. A belief that vagueness is inessential and therefore unimportant may comfort those who ignore the phenomenon. But their complacency is unjustified. Even if we could reduce the vagueness in our language (as science is often described as striving to do by producing sharper definitions, and as legal processes can accomplish via appeal to precedents), our efforts could not in practice eliminate it entirely. (Russell himself stresses the persistent vagueness in scientific terms, p. 86; and it is clear that the legal process could never reach absolute precision either.) Moreover, in natural language vague predicates are ubiquitous, and this alone motivates study of the phenomenon irrespective of whether there could be usable languages entirely free of vagueness. Even if ‘heap’ could be replaced by some term ‘heap★’ with perfectly sharp boundaries and for which no sorites paradox would arise, the paradox facing our actual vague term would remain.² And everyday reasoning takes place in vague language, so no account of good ordinary reasoning can ignore vagueness.

² See Carnap 1950, chapter 1, Haack 1974, chapter 6 and Quine 1981 on the replacement of vague expressions by precise ones, and see Grim 1982 for some difficulties facing the idea. Certain predicates frequently prompt the response that there *is* in fact a sharp boundary for their strict application, though we use them more loosely – in par-

In the next section I shall discuss the variety of vague expressions – a variety which is not brought out by the general form of arguments for tolerance. First, I clarify the phenomenon by mentioning three things that vagueness in our sense (probably) is not.

(a) The remark ‘Someone said something’ is naturally described as vague (who said what?). Similarly, ‘X is an integer greater than thirty’ is an unhelpfully vague hint about the value of X. Vagueness in this sense is underspecificity, a matter of being less than adequately informative for the purposes in hand. This seems to have nothing to do with borderline cases or with the lack of sharp boundaries: ‘is an integer greater than thirty’ has sharp boundaries, has no borderline cases, and is not susceptible to sorites paradoxes. And it is not because of any possibility of borderline people or borderline cases of saying something that ‘someone said something’ counts as vague in the alternative sense. I shall ignore the idea of vagueness as underspecificity: in philosophical contexts, ‘vague’ has come to be reserved for the phenomenon I have described.

(b) Vagueness must not be straightforwardly identified with paradigm context-dependence (i.e. having a different extension in different contexts), even though many terms have both features (e.g. ‘tall’). Fix on a context which can be made as definite as you like (in particular, choose a specific comparison class, e.g. current professional American basketball players): ‘tall’ will remain vague, with borderline cases and fuzzy boundaries, and the sorites paradox will retain its force. This indicates that we are unlikely to understand vagueness or solve the paradox by concentrating on context-dependence.³

(c) We can also distinguish vagueness from ambiguity. Certainly, terms can be ambiguous *and* vague: ‘bank’ for example has two quite different main senses (concerning financial institutions or river edges), both of which are vague. But it is natural to suppose that ‘tadpole’ has a univocal sense, though that sense does not determine a sharp, well-defined extension. Certain theories, however, do

ticular, strictly no one is bald unless they have absolutely no hair (see e.g. Sperber and Wilson 1986). But even if this line is viable in some cases, it is hopeless for the majority of vague predicates. E.g. should someone count as ‘tall’ only if they are as tall as possible? How about ‘quite tall’? Or ‘very hairy’? And where is the strict boundary of ‘chair’?

³ There have, however, been some attempts at this type of solution to the sorites paradox using, for example, more elaborate notions of the context of a subject’s judgement (see e.g. Raffman 1994).

attempt to close the gap between vagueness and a form of ambiguity (see chapter 7, §1).

2. TYPE OF VAGUE EXPRESSIONS

So far, I have focused on a single dimension of variation associated with each vague predicate, such as height for ‘tall’ and number of grains for ‘heap’. But many vague predicates are *multi-dimensional*: several different dimensions of variation are involved in determining their applicability. The applicability of ‘big’, used to describe people, depends on both height and volume; and even whether something counts as a ‘heap’ depends not only on the number of grains but also on their arrangement. And with ‘nice’, for example, there is not even a clear-cut set of dimensions determining the applicability of the predicate: it is a vague matter which factors are relevant and the dimensions blend into one another.

The three central features of vague predicates are shared by multi-dimensional ones. There are, for example, borderline nice people: indeed, some are borderline *because* of the multi-dimensionality of ‘nice’, by scoring well in some relevant respects but not in others. Next consider whether multi-dimensional predicates may lack sharp boundaries. In the one-dimensional case, F has a sharp boundary (or sharp boundaries) if possible candidates for it can be ordered with a point (or points) marking the boundary of F ’s extension, so that everything that falls on one side of the point (or between the points) is F and nothing else is F . For a multi-dimensional predicate, there may be no uniquely appropriate ordering of possible candidates on which to place putative boundary-marking points. (For instance, there is no definite ordering of people where each is bigger than the previous one; in particular, if ordered by height, volume is ignored, and vice versa.) Rather, for a sharply bounded two-dimensional predicate the candidates would be more perspicuously set out in a two-dimensional space in which a boundary could be drawn, where the two-dimensional region enclosed by the boundary contains all and only instances of the predicate. With a *vague* two-dimensional predicate no such sharp boundary can be drawn. Similarly, for a sharply bounded predicate with a clear-cut set of n dimensions, the boundary would enclose an n -dimensional region containing all of its instances; and vague predicates will lack such a sharp

boundary.⁴ When there is no clear-cut set of dimensions – for ‘nice’, for example – this model of boundary-drawing is not so easily applied: it is then not possible to construct a suitable arrangement of candidates on which to try to draw a boundary of the required sort. But this, I claim, is distinctive of the vagueness of such predicates: they have no sharp boundary, but nor do they have a fuzzy boundary in the sense of a rough boundary-area of a representative space. ‘Nice’ is so vague that it cannot even be associated with a neat array of candidate dimensions, let alone pick out a precise area of such an array.

Finally, multi-dimensional vague predicates are susceptible to sorites paradoxes. We can construct a sorites series for ‘heap’ by focusing on the number of grains and minimising the difference in the arrangement of grains between consecutive members. And for ‘nice’ we could take generosity and consider a series of people differing gradually in this respect, starting with a very mean person and ending with a very generous one, where, for example, other features relevant to being nice are kept as constant as possible through the series.

Next, I shall argue that comparatives as well as monadic predicates can be vague. This has been insufficiently recognised and is sometimes denied. Cooper 1995, for example, seeks to give an account of vagueness by explaining how vague monadic predicates depend on comparatives, taking as a starting point the claim that ‘classifiers in their grammatically positive form [e.g. “large”] are vague, while comparatives are not’ (p. 246). With a precise comparative, ‘*F*-er than’, for any pair of things *x* and *y*, either *x* is *F*-er than *y*, *y* is *F*-er than *x*, or they are equally *F*. This will be the case if there is a determinate ordering of candidates for *F*-ness (allowing ties). For example, there is a one-dimensional ordering of the natural numbers relating to the comparative ‘is a smaller number than’, and there are no borderline cases of this comparative, which is paradigmatically precise. Since ‘is a small number’ is a vague predicate, this shows how vague positive forms can have precise comparatives. It may seem that

⁴ Could there be a single, determinate way of balancing the various dimensions of a multi-dimensional predicate that *does* yield a unique ordering? Perhaps, but this will usually not be the case, and when it is, it may then be appropriate to treat the predicate as one-dimensional, even if the ‘dimension’ is not a natural one. Further discussion of this point would need a clearer definition of ‘dimension’, but this is not important for our purposes.

‘older than’ also gives rise to an ordering according to the single dimension of age, and hence that ‘older than’ must be precise. But, in fact, there could be borderline instances of the comparative due to indeterminacy over exactly what should count as the instant of someone’s birth and so whether it is before or after the birth of someone else. And such instances illustrate that there is *not*, in fact, an unproblematic ordering of *people* for ‘older than’, even though there is a total ordering of *ages*, on which some people cannot be exactly placed. Similarly, though there is a single dimension of height, people cannot always be exactly placed on it and assigned an exact height. For what exactly should count as the top of one’s head? Consequently there may also be borderline instances of ‘taller than’.

Comparatives associated with multi-dimensional predicates – for example ‘nicer than’ and ‘more intelligent than’ – are typically vague. They have borderline cases: pairs of people about whom there is no fact of the matter about who is nicer/more intelligent, or whether they are equally nice/intelligent. This is particularly common when comparing people who are nice/intelligent in different ways. There are, however, still clear cases of the comparative in addition to borderline cases – it is not that people are *never* comparable in respect to niceness – thus the vague ‘nicer than’, like ‘nice’ itself, has clear positive, clear negative and borderline cases.

Can comparatives also lack sharp boundaries? Talk of boundaries, whether sharp or fuzzy, is much less natural for comparatives than for monadic predicates. But we might envisage precise comparatives for which we could systematically set out ordered pairs of things, $\langle x, y \rangle$ and draw a sharp boundary around those for which it is true that x is F -er than y . For example, if F has a single dimension then we could set out pairs in a two-dimensional array, where the x co-ordinate of a pair is determined by the location along the dimension of the first of the pair, and the y co-ordinate by that of the second. The boundary line could then be drawn along the diagonal at $x = y$, where pairs falling beneath the diagonal are definitely true instances of the comparative ‘ x is F -er than y ’, and those on or above are definitely false. But for many comparatives, including ‘nicer than’, there could not be such an arrangement and this gives a sense in which those comparatives lack sharp boundaries.

Another possible sense in which comparatives may lack sharp boundaries is the following. Take the comparative ‘redder than’ and

choose a purplish-red patch of colour, a . Then consider a series of orangeish-red patches, x_i , where x_{i+1} is redder than x_i . It could be definitely true that a is redder than x_0 (which is nearly orange), definitely not true that a is redder than x_{100} , where not only are there borderline cases of ' a is redder than x_i ' between them, but there is no point along the series of x_i at which it suddenly stops being the case that a is redder than x_i . So, certain comparatives have borderline cases and exhibit several features akin to the lack of sharp boundaries: they should certainly be classified as vague.

Having discussed vague monadic predicates and vague comparatives, I shall briefly mention some other kinds of vague expressions. First, there can be other vague dyadic relational expressions. For example, 'is a friend of' has pairs that are borderline cases. Adverbs like 'quickly', quantifiers like 'many' and modifiers like 'very' are also vague. And, just as comparatives can be vague, particularly when related to a multi-dimensional positive, so can superlatives. 'Nicest' and 'most intelligent' have vague conditions of application: among a group of people it may be a vague matter, or indeterminate, who is the nicest or the most intelligent. And vague superlatives provide one way in which to construct vague singular terms such as 'the nicest man' or 'the grandest mountain in Scotland', where there is no fact of the matter as to which man or mountain the terms pick out. Terms with plural reference like 'the high mountains of Scotland' can equally be vague.

A theory of vagueness should have the resources to accommodate all the different types of vague expression. And, for example, we should reject an account of vagueness that was obliged to deny the above illustrated features of certain comparatives in order to construct its own account of vague monadic predicates. (See chapter 5, §2 about this constraint in connection with degree theories.) The typical focus on monadic predicates need not be mistaken, however. Perhaps, as Fine suggests, all vagueness is reducible to predicate vagueness (1975, p. 267), though such a claim needs supporting arguments. Alternatively, vagueness might manifest itself in different ways in different kinds of expression, and this could require taking those different expression-types in turn and having different criteria of vagueness for comparatives and monadic predicates. Another possibility is to treat complete sentences as the primary bearers of vagueness, perhaps in their possession of a non-classical truth-value.

This approach would avoid certain tricky questions about whether the vagueness of a particular sentence is ‘due to’ a given expression. For example, in a case where it is indeterminate exactly what moment *a* was born and whether it was before the birth of *b*, we would avoid the question whether this shows ‘older than’ to be vague, or whether the indeterminacy should be put down to vagueness in *a* itself. Provided one can still make sense of a typical attribution of vagueness to some element of a sentence in the uncontroversial cases, I suggest that this strategy is an appealing one.

3. VAGUENESS IN THE WORLD?

Is it only linguistic items – words or phrases – that can be vague? Surely not: thoughts and beliefs are among the mental items which share the central characteristics of vagueness; other controversial cases include perceptions. What about the world itself: could the world be vague as well as our descriptions of it? Can there be vague objects? Or vague properties (the ontic correlates of predicates)? Consider Ben Nevis: any sharp spatio-temporal boundaries drawn around the mountain would be arbitrarily placed, and would not reflect a natural boundary. So it may seem that Ben Nevis has fuzzy boundaries, and so, given the common view that a vague object is an object with fuzzy, spatio-temporal boundaries, that it is a vague object. (See e.g. Parsons 1987, Tye 1990 and Zemach 1991 for arguments that there are vague objects.) But there are, of course, other contending descriptions of the situation here. For example, perhaps the only objects we should admit into our ontology are precise/sharp although we fail to pick out a single one of them with our (vague) name ‘Ben Nevis’. It would then be at the level of our representations of the world that vagueness came in. (See chapter 7, §1 on an indeterminate reference view.)⁵

My concern is with linguistic vagueness and I shall generally ignore ontic vagueness. This would be a mistake if a theory of linguistic vagueness had to rely on ontic vagueness. But that would be surprising since it seems at least possible to have vague language in a non-vague world. In particular, even if all objects, properties and

⁵ The most discussed strand of the ontic vagueness debate focuses on Evans’s formal argument which aims to establish a negative answer to his question ‘Can there be vague objects?’ (1978; see Keefe and Smith 1997b, §5 for an overview of the debate).

facts were precise, we would still have reason, for everyday purposes, to use a vague expression such as ‘tall’, which would still have borderline cases (even if those cases could also be described in non-vague terms involving precise heights etc.). Similarly, in a precise world we would still use vague singular terms, perhaps to pick out various large collections of precise fundamental particulars (e.g. as clouds or mountains) where the boundaries of those collections are left fuzzy. So it seems that language could still be vague if the concrete world were precise.⁶

The theories of vagueness of this book are theories of linguistic vagueness and in the next section I briefly introduce them.

4. THEORIES OF VAGUENESS

The candidate theories of vagueness can be systematically surveyed by considering how they address two central tasks. The first is to identify the logic and semantics for a vague language – a task bound up with providing an account of borderline cases and of fuzzy boundaries. The second task is that of addressing the sorites paradox.

(i) *The logic and semantics of vagueness*

The simplest approach is to retain classical logic and semantics. Borderline case predications *are* either true or false after all, though we do not and cannot know which. Similarly, despite appearances, vague predicates have well-defined extensions: there is a sharp boundary between the tall people and the rest, and between the red shades of the spectrum and the other colours. As chapter 3 will describe, the *epistemic view* takes this line and accounts for vagueness in terms of our ignorance – for example, ignorance of where the sharp boundaries to our vague predicates lie. And a *pragmatic account* of vagueness also seeks to avoid challenging classical logic and semantics, but this time by accounting for vagueness in terms of pragmatic relations between speakers and their language: see chapter 6.

⁶ These are only *prima facie* reasons for not approaching linguistic vagueness via ontic vagueness: a tighter case would require clarification of what vagueness in the world would be. They also do not seem to bear on the question whether there can be vague sets, which might also be counted as a form of ontic vagueness. Tye, for example, believes that there are vague sets and maintains that they are crucial to his own theory of the linguistic phenomena (see Tye 1990).

If we do not retain classical logic and semantics, we can say instead that when a is a borderline case of F , the truth-value of ' a is F ' is, as Machina puts it, 'in some way peculiar, or indeterminate or lacking entirely' (1976, p. 48). This generates a number of non-classical options.

Note that a borderline case of the predicate F is equally a borderline case of not- F : it is unclear *whether or not* the candidate is F . This symmetry prevents us from simply counting a borderline F as not- F . But there are several ways of respecting this symmetry. Some take the line that a predication in a borderline case is both true *and* false: there is a truth-value glut. This can be formalised within the context of a paraconsistent logic – a logic that admits true contradictions (see Hyde 1997 and chapter 7, §7 for discussion of that view).

A more popular position is to admit truth-value *gaps*: borderline predications are neither true nor false. One elegant development is *supervaluationism*. The basic idea is that a proposition involving the vague predicate 'tall', for example, is true (false) if it comes out true (false) on all the ways in which we can make 'tall' precise (ways, that is, which preserve the truth-values of uncontentionally true or false cases of ' a is tall'). A borderline case, 'Tek is tall', will be neither true nor false, for it is true on some ways of making 'tall' precise and false on others. But a classical tautology like 'either Tek is tall or he is not tall' *will* still come out true because wherever a sharp boundary for 'tall' is drawn, that compound sentence will come out true. In this way, the supervaluationist adopts a non-classical *semantics* while aiming to minimise divergence from classical *logic*. A theory of this type will be defended in chapters 7 and 8.

Rather than holding that predications in borderline cases lack a truth-value, another option is to hold that they have a third value – 'neutral', 'indeterminate' or 'indefinite' – leading to a three-valued logic (see chapter 4). Alternatively, *degree theories* countenance degrees of truth, introducing a whole spectrum of truth-values from 0 to 1, with complete falsity as degree 0 and complete truth as degree 1. Borderline cases each take some value between 0 and 1, with ' x is red' gradually increasing in truth-value as we move along the colour spectrum from orange to red. This calls for an infinite-valued logic or a so-called 'fuzzy logic', and there have been a variety of different versions (see chapter 4).

So far the sketched positions at least agree that there is *some* positive

account to be given of the logic and semantics of vagueness. Other writers have taken a more pessimistic line. In particular, Russell claims that logic assumes precision, and since natural language is not precise it cannot be in the province of logic at all (1923, pp. 88–9). If such a ‘no logic’ thesis requires wholesale rejection of reasoning with vague predicates – and hence of most reasoning in natural language – it is absurdly extreme. And arguments involving vague predicates are clearly not all on a par. For example, ‘anyone with less than 500 hairs on his head is bald; Fred has less than 500 hairs on his head; therefore Fred is bald’ is an unproblematically good argument (from Cargile 1969, pp. 196–7). And, similarly, there are other ways of arguing with vague predicates that should certainly be rejected. Some account is needed of inferences that are acceptable and others that fail, and to search for systematic principles capturing this is to seek elements of a logic of vague language. So, I take the pessimism of the no-logic approach to be a very last resort, and in this book I concentrate on more positive approaches.

Focusing on the question how borderline case predication should be classified, we seem to have exhausted the possibilities. They may be true or false, or have no truth value at all (in particular, being neither true nor false), or be both true and false, or have a non-classical value from some range of values. When it comes to surveying solutions to the sorites paradox, however, there may additionally be alternatives that do not provide a theory of vagueness and perhaps do not answer the question how borderline cases are to be classified. I concentrate on those which do fit into a theory of vagueness.

(ii) *The sorites paradox*

A paradigm sorites set-up for the predicate F is a sequence of objects x_i , such that the two premises

- (1) Fx_1
- (2) For all i , if Fx_i then Fx_{i+1}

both appear true, but, for some suitably large n , the putative conclusion

- (3) Fx_n

seems false. For example, in the case of ‘tall’, the x_i might be the series of men described earlier, each a hundredth of an inch shorter

than the previous one and where x_1 is seven feet tall. (1) ' x_1 is tall' is then true; and so, it seems, is the inductive premise, (2) 'for all i , if x_i is tall, so is x_{i+1} '. But it is surely false that (3) x_{3000} – who is only 4 feet 6 inches – is tall.

A second form of sorites paradox can be constructed when, instead of the quantified inductive premise (2), we start with a collection of particular conditional premises, $(2C_i)$, each of the form 'if Fx_i then Fx_{i+1} '. For example,

(2C₁) if x_1 is tall, so is x_2

(2C₂) if x_2 is tall, so is x_3

and so on. And the use of conditionals is not essential: we can take a sequence of premises of the form $\neg(Fx_i \ \& \ \neg Fx_{i+1})$ – a formulation that goes back at least to Diogenes Laertius (see Long and Sedley 1987, p. 222). Alternatively, (2) could be replaced by a quantification over the negated conjunctions of that form.

As well as needing to solve the paradox, we must assess that general form of argument because it is used both in philosophical arguments outside the discussion of vagueness (e.g. with the story of the ship of Theseus) and in various more everyday debates (the slippery slope arguments mentioned in §1).⁷

Responses to a sorites paradox can be divided into four types. We can:

(a) deny the validity of the argument, refusing to grant that the conclusion follows from the given premises; or

(b) question the strict truth of the general inductive premise (2) or of at least one of the conditionals $(2C_i)$; or

(c) accept the validity of the argument and the truth of its inductive premise (or of all the conditional premises) but contest the supposed truth of premise (1) or the supposed falsity of the conclusion (3); or

(d) grant that there are compelling reasons both to take the

⁷ As a further example of the former, consider Kirk 1986 (pp. 217ff). Regarding Quine's thesis about the indeterminacy of translation, Kirk uses an argument with the form of the quantificational version of the paradox to argue that there can be no indeterminacy of translation because, first, there would be no indeterminacy in translating between the languages of infants each of whom is at an early stage of language-acquisition and, second, if there is no indeterminacy at one step of acquisition then there is none at the next. He presents his argument as using mathematical induction but does not ask whether its employment of vague predicates casts doubt on that mode of argument.

argument form as valid, and to accept the premises and deny the conclusion, concluding that this demonstrates the incoherence of the predicate in question.

I shall briefly survey these in turn, ignoring here the question whether we should expect a uniform solution to all sorites paradoxes whatever their form and whatever predicate is involved. (Wright 1987 argues that different responses could be required depending on the reasons that support the inductive premise.) Any response must explain away apparent difficulties with accepting the selected solution; for example, if the main premise is denied, it must be explained why that premise is so plausible. More generally, a theory should account for the persuasiveness of the paradox as a paradox and should explain how this is compatible with the fact that we are never, or very rarely, actually led into contradiction.

(a) Denying the validity of the sorites argument seems to require giving up absolutely fundamental rules of inference. This can be seen most clearly when the argument takes the second form involving a series of conditionals, the $(2C_i)$. The only rule of inference needed for this argument is modus ponens. Dummett argues that this rule cannot be given up, as it is constitutive of the meaning of 'if' that modus ponens is valid (1975, p. 306). To derive the conclusion in the first form of sorites, we only need universal instantiation in addition to modus ponens; but, as Dummett again argues, universal instantiation seems too central to the meaning of 'all' to be reasonably challenged (1975, p. 306). I agree on both points and shall not pursue the matter further here.

There is, however, a different way of rejecting the validity of the many-conditionals form of the sorites. It might be suggested that even though each step is acceptable on its own, chaining too many steps does not guarantee the preservation of truth if what counts as preserving truth is itself a vague matter. (And then the first form of sorites could perhaps be rejected on the grounds that it is in effect short hand for a multi-conditional argument.) As Dummett again notes, this is to deny the transitivity of validity, which would be another drastic move, given that chaining inferences is normally taken to be essential to the very enterprise of proof.⁸

⁸ But see Parikh 1983. In my chapter 4, §7 the possibility is briefly entertained.

Rather than questioning particular inference rules or the ways they can be combined, Russell's global rejection of logic for vague natural language leads him to dismiss 'the old puzzle about the man who went bald', simply on the grounds that 'bald' is vague (1923, p. 85). The sorites arguments, on his view, cannot be valid because, containing vague expressions, they are just not the kind of thing that can be valid or invalid.

(b) If we take a formulation of the paradox that uses negated conjunctions (or assume that 'if' is captured by the material conditional), then within a classical framework denying the quantified inductive premise or one of its instances commits us to there being an i such that ' Fx_i and not- Fx_{i+1} ' is true. This implies the existence of sharp boundaries and the epistemic theorist, who takes this line, will explain why vague predicates appear not to draw sharp boundaries by reference to our ignorance (see chapter 3).

In a non-classical framework there is a wide variety of ways of developing option (b), and it is not clear or uncontroversial which of these entail a commitment to sharp boundaries. For example, the supervaluationist holds that the generalised premise (2) 'for all i , if Fx_i then Fx_{i+1} ' is false: for each F^\star which constitutes a way of making F precise, there will be *some* x_i or other which is the last F^\star and is followed by an x_{i+1} which is not- F^\star . But since there is no particular i for which ' Fx_i and not- Fx_{i+1} ' is true – i.e. true however F is made precise – supervaluationists claim that their denial of (2) does not mean accepting that F is sharply bounded (see chapter 7). And other non-classical frameworks may allow that (2) is not true, while not accepting that it is *false*. Tye 1994, for example, maintains that the inductive premise and its negation both take his intermediate truth-value, 'indefinite'.

Degree theorists offer another non-classical version of option (b): they can deny that the premises are strictly true while maintaining that they are *nearly* true. The essence of their account is to hold that the predications Fx_i take degrees of truth that encompass a gradually decreasing series from complete truth (degree 1) to complete falsity (degree 0). There is never a substantial drop in degree of truth between consecutive Fx_i ; so, given a natural interpretation of the conditional, the particular premises 'if Fx_i then Fx_{i+1} ' can each come out at least very nearly true, though some are not completely true. If the sorites argument based on many conditionals is to count as strictly valid, then

an account of validity is needed that allows a valid argument to have nearly true premises but a false conclusion. But with some degree-theoretic accounts of validity, the sorites fails to be valid – thus a degree theorist can combine responses (a) and (b) (see chapter 4, §7).

Intuitionistic logic opens up the possibility of another non-classical position that can respond to the sorites by denying the inductive premise (2), while not accepting the classical equivalent of this denial, $(\exists x_i)(Fx_i \ \& \ \neg Fx_{i+1})$, which is the unwanted assertion of sharp boundaries. Putnam 1983 suggests this strategy. But critics have shown that with various reasonable additional assumptions, other versions of sorites arguments still lead to paradox. In particular, if, as might be expected, you adopt intuitionistic semantics as well as intuitionistic logic, paradoxes recur (see Read and Wright 1985). And Williamson 1996 shows that combining Putnam's approach to vagueness with his epistemological conception of truth still faces paradox. (See also Chambers 1998, who argues that, given Putnam's own view on what would make for vagueness, paradox again emerges.) The bulk of the criticisms point to the conclusion that there is no sustainable account of vagueness that emerges from rejecting classical logic in favour of intuitionistic logic.

(c) Take the sorites (H+) with the premises 'one grain of sand is not a heap' and 'adding a single grain to a non-heap will not turn it into a heap'. If we accept these premises and the validity of the argument, it follows that we will never get a heap, no matter how many grains are piled up: so there are no heaps. Similarly, sorites paradoxes for 'bald', 'tall' and 'person' could be taken to show that there are no bald people, no tall people and indeed no people at all. Unger bites the bullet and takes this nihilistic line, summarised in the title of one of his papers: 'There are no ordinary things' (Unger 1979; see also Wheeler 1975, 1979 and Heller 1988).

The thesis, put in linguistic terms, is that all vague predicates lack serious application, i.e. they apply either to nothing ('is a heap') or to everything ('is not a heap'). Classical logic can be retained in its entirety, but sharp boundaries are avoided by denying that vague predicates succeed in drawing *any* boundaries, fuzzy or otherwise. There will be no borderline cases: for any vague F , everything is F or everything is not- F , and thus nothing is borderline F .⁹

⁹ See Williamson 1994, chapter 6, for a sustained attack on various forms of nihilism. For example, he shows how the nihilist cannot state or argue for his own position on

The response of accepting the conclusion of every sorites paradox cannot be consistently sustained. For in addition to (H+), there is the argument (H−) with the premises ‘ten thousand grains make a heap’ and ‘removing one grain from a heap still leaves a heap’, leading to the conclusion that a single grain of sand is a heap, which is incompatible with the conclusion of (H+). Such reversibility is typical; given a sorites series of items, the argument can be run either way through them. Unger’s response to (H−) would be to deny the initial premise: there are no heaps – as (H+) supposedly shows us – so it is not true that ten thousand grains make a heap. Systematic grounds would then be needed to enable us to decide which of a pair of sorites paradoxes is sound (e.g. why there are no heaps rather than everything being a heap).

Unger is driven to such an extreme position by the strength of the arguments in support of the inductive premises of sorites paradoxes. If our words determined sharp boundaries, Unger claims, our understanding of them would be a miracle of conceptual comprehension (1979, p. 126). The inductive premise, guaranteeing this lack of sharp boundaries, reflects a semantic rule central to the meaning of the vague *F*. But, we should ask Unger, can the tolerance principle expressed in the inductive premise for ‘tall’ really be more certain than the truth of the simple predication of ‘tall’ to a seven-foot man? Is it plausible to suppose that the expression ‘tall’ is meaningful and consistent but that there could not be anything tall, when learning the term typically involves ostension and hence confrontation with alleged examples? A different miracle of conceptual comprehension would be needed then to explain how we can understand that meaning and, in general, how we can use such empty predicates successfully to communicate anything at all. It may be more plausible to suppose that if there are any rules governing the application of ‘tall’, then, in addition to tolerance rules, there are ones dictating that ‘tall’ applies to various paradigmatic cases and does not apply to various paradigmatically short people. Sorites paradoxes could then demonstrate the inconsistency of such a set of rules, and this is option (d).

Responses (c) and (d) are not always clearly distinguished. Writers

his own terms (e.g. the expressions he tries to use must count as incoherent). My discussion of methodological matters in chapter 2 will suggest that a swifter rejection of the position is warranted anyway.

like Unger are primarily concerned with drawing ontological conclusions. It is enough for them to emphasise the tolerance of a predicate like 'tall' which already guarantees, they claim, that the world contains nothing that strictly answers to that description: they are not so concerned to examine what further rules might govern the predicate and perhaps render it incoherent. But other writers, for example Dummett, explore these conceptual questions.

(d) Having argued in detail against alternative responses to the paradox, Dummett 1975 maintains that there is no choice but to accept that a sorites paradox for *F* exemplifies an undeniably valid form of argument from what the semantic rules for *F* dictate to be true premises to what they dictate to be a false conclusion. The paradoxes thus reveal the incoherence of the rules governing vague terms: by simply following those rules, speakers could be led to contradict themselves. This inconsistency means that there can be no coherent logic governing vague language.¹⁰

Once (d)-theorists have concluded that vague predicates are incoherent, they may agree with Russell that such predicates cannot appear in valid arguments. So option (d) can be developed in such a way that makes it compatible with option (a), though this route to the denial of validity is very different from Russell's. (Being outside the scope of logic need not make for incoherence.)

The acceptance of such pervasive inconsistency is highly undesirable and such pessimism is premature; and it is even by Dummett's own lights a pessimistic response to the paradox, adopted as a last resort rather than as a positive treatment of the paradox that stands as competitor to any other promising alternatives. Communication using vague language is overwhelmingly successful and we are never in practice driven to incoherence (a point stressed by Wright, e.g. 1987, p. 236). And even when shown the sorites paradox, we are rarely inclined to revise our initial judgement of the last member of the series. It looks unlikely that the success and coherence in our practice is owed to our grasp of inconsistent rules. A defence of some version of option (a) or (b) would provide an attractive way of

¹⁰ See also Rolf 1981, 1984. Horgan 1994, 1998 advocates a different type of the inconsistency view. He agrees that sorites paradoxes (and other related arguments) demonstrate logical incoherence, but considers that incoherence to be *tempered* or *insulated*, so that it does not infect the whole language and allows us to use the language successfully despite the incoherence (see chapter 8, §2).